

SmarTxT: A Natural Language Processing Approach for Efficient Vehicle Defect Investigation ¹

Jack Francis

Davidson Technologies Inc.
Huntsville, AL

Kim Cates

KBR, Inc.
Lexington Park, MD
kim.cates@us.kbr.com

Gianluigi Caldiera, Corresponding Author

KBR, Inc.
Fulton, MD
gianluigi.caldiera@us.kbr.com

January 2022

ABSTRACT

The investigation of vehicle defects, which are generally led by the National Highway Transportation Safety Administration (NHTSA) in the United States, is critical to the continued trust of the general public in the safety of vehicles. NHTSA routinely receives millions of reports of potential defects, complaints, recalls, and manufacturer communications, which may provide evidence of a new vehicle defect. However, the large quantity and text-based communication make efficiently identifying defect trends difficult for analysts. To accelerate the investigation of defect reports, we introduce a Natural Language Processing (NLP) application that identifies key topics and similar defect reports to assist analysts and investigators. Further, our application is built to provide users with a web interface for interacting with the NLP models. The integration of NLP with current NHTSA datasets provides a method for quickly identifying defect trends in large text-based datasets. To demonstrate the effectiveness of our method, we apply our approach to two publicly available NHTSA datasets, namely the Technical Service Bulletins and Recalls dataset.

Keywords: Natural Language Processing, Topic Extraction, Document Similarity, Transformers, SmarTxT

This work was funded by KBR Science and Space and would not have been possible without the support of KBR and support from Jim King, Kelly Kim, Rohit Mital, and Todd May.

1. INTRODUCTION

Like in many other domains, the size and frequency of data has increased exponentially in recent years in transportation-related fields. This drastic increase in data availability has led to artificial intelligence and machine learning (AI/ML) techniques being frequently used for a wide range of transportation applications. This effect has led to the number of AI/ML papers to the Transportation Research Board Annual Meeting to skyrocket from 5 in 2015 to 162 in 2020 [1]. Most of the AI/ML work is concentrated on navigation applications, advanced driver assistance, and smart traffic systems [2, 3, 4, 5, 6, 5, 7]. A common technology used for these applications, deep learning, is often used to detect defects in the manufacture of vehicles and vehicle components [7, 8]. In these applications, researchers use convolutional neural networks, recurrent neural networks, and depth residual neural networks, to name a few, to extract features from images and map them to categories of defects. The investigation and resolution of defects in vehicles are critical to ensure public safety and confidence.

Two key challenges arise when an analyst, from the National Highway Traffic Safety Administration (NHTSA) or a safety advocacy organization, like the Insurance Institute for Highway Safety (IIHS), wants to investigate and address vehicle defects. First, defect data is generally recorded in text with associated metadata, such as the make, model, and year of the vehicle. The text data makes performing AI/ML approaches difficult, because most classical ML algorithms do not work with text data directly [9]. Because of this, an analyst must read each defect record line by line to understand the defect and identify patterns among defects. For example, analysts in NHTSA's Office of Defect Investigation (ODI) identify defect issues through manual evaluation of consumer complaints, equipment defects, and injury reports. The results of these manual preliminary investigations are summarized and stored in the Artemis database [10, 11]. Although historical investigations are stored in Artemis, there is currently no method for identifying the similarity of a current investigation to any past historical investigations. This is a key challenge for ODI analysts, who often examine defects that occur on components used in multiple makes and models. Second, the data size is exceedingly large. NHTSA publicly posts a variety of datasets relating to defects, including consumer complaints, recalls, and manufacturer communications. These datasets have millions of entries, which makes the analyst's job of manually reading defect reports line by line to identify patterns essentially infeasible. Natural language processing (NLP) techniques, that will be introduced in this paper, help the analyst in dealing with this problem. NLP is a subfield of machine learning that aims to train computers to understand spoken words and text like humans and is further described in Section 2.

In this paper, we propose an application, namely SmarTxT, that not only efficiently processes the large text datasets of defects using Natural Language Processing (NLP) but also provides a user-friendly interface to assist analysts in the investigation and resolution of defects. To enhance the analyst's investigation, we implement two NLP approaches, namely topic modeling and document similarity. Topic modeling identifies latent topics for a given dataset, which provides an analyst with the major trends for a data selection. Document similarity provides analysts with similar defect records. To implement document similarity, we use BERT (Bidirectional Encoder

Representations from Transformers) to embed defect records and then use cosine similarity to identify similarity among defect records. To demonstrate the effectiveness of our NLP approaches, we apply these methods to the NHTSA Technical Service Bulletins (TSBS) dataset as a case study. These methods are unsupervised, so the model results are validated by a human expert. Finally, we built the application using Flask and Python to provide an accessible approach for analysts to interact with our modeling approaches.

Our proposed application built on NLP provides major benefits to analysts in defect investigations:

1. Our NLP approach provides a fast, accurate method for investigating millions of defect records. These approaches significantly reduce the number of hours analysts must spend on reading individual defect records. Further, the immediate aggregation of similar documents may lead to the identification of new patterns and trends that would have previously been missed.
2. SmarTxT provides an easy and accessible interface for analysts to understand the results from our NLP models. SmarTxT is built so that NLP model understanding is not required for analysts, which greatly increases the number of analysts that can benefit from this tool. It is a tool designed for what Gartner Group has described as the “citizen data scientist”.

The remainder of this paper is structured as follows. Section 2 provides a description of the data and our preprocessing approach to convert the data into a format appropriate for NLP. The NLP approaches, namely topic modeling and document similarity, are described in Section 3. Section 4 provides an overview of the SmarTxT app. Section 5 details our validation approach for assessing model quality in determining similar documents. Finally, Section 6 concludes our study with conclusions and future work.

2. LITERATURE REVIEW

NLP is a subfield of machine learning that aims to understand the contents of documents, where a document in the NLP setting is any collection of text, which could include words, sentences, paragraphs, and/or papers [12, 13]. The key challenge in NLP is associating documents with a lower-level embedding that captures the semantic similarity and context between two documents. There have been multiple methods proposed for this problem, including Word2Vec [13], Doc2Vec [14], TF-IDF (term frequency-inverse document frequency) [15], topic modeling [16], and more recently transformer models [17]. In each of these methods, the objective is to take in a corpus of documents and assign a lower-dimension numerical embedding to capture the relationship between documents. Because documents are defined as a collection of documents, there are varying levels of abstraction that each of these models attempts to embed. For example, Word2Vec operates on individual words, Doc2Vec and TF-IDF operate on paragraphs, and transformers can be used for either use case depending on the model selected. Some of the widely used transformer models for NLP are BERT [18], ELMo [19], and GPT-3 [20]. These transformer models are trained on large corpora and distributed with the learned weights from this training procedure. For example, BERT was trained using the BooksCorpus (800M words) and English Wikipedia (2,500M words) [18]. A major benefit of these large pretrained transformer models is the relative ease of applying

transfer learning for new domains [21]. By freezing all weights except the top layer and then retraining the model on domain-specific corpuses the models can be quickly adapted for new domains. Examples of domain-specific retraining include BioBERT [22] and LegalBERT [23].

Due to the large success of NLP over the recent years, these methods have been applied in a wide range of applications. Of particular interest for SmarTxT are applications that leverage topic modeling, document similarity, and transformer models. Topic modeling is an NLP technique for identifying latent topics from a set of documents and is described in further detail in Section 4.1. Topic modeling has been used in a wide variety of applications including engineering, bioinformatics, and automated vehicles [24, 25, 26, 27]. Document similarity is a natural extension of embeddings, where we identify the most similar document to a selected reference document. A common use case for document similarity is to return contextually relevant information for search queries [28]. Additional applications of document similarity include providing semantic search capability for aviation maintenance records [29] and searching clinical notes in health centers [30].

Topic modeling and document similarity can enhance the current data review and investigation of defects by NHTSA. First, the identification of latent topics in maintenance records provides an overview of trending defects. The identification of topics allows NHTSA ODI analysts to evaluate trending defects and then compare against previous defect investigations to identify long-term defect trends. Currently, this is not feasible due to the large amount of data and the dynamic SOPs [11]. Further, topic modeling allows for cross-make and cross-model analyses to be completed. By combining all documents from various makes and models, analysts can identify defect trends among parts shared between multiple automakers. Second, document similarity provides a semantic, contextual search for analysts to identify similar defects. Defect records are not written in a consistent format, which leads to further challenges in comparing cross-make and cross-model defect reports. With document similarity, NHTSA ODI analysts would have the ability to input full defect descriptions or key phrases to query semantically similar records within the Artemis database. Our solution, SmarTxT, proposes to leverage NLP techniques for vehicle defect investigation, which to the author's knowledge has not previously been completed.

3. DATASET DESCRIPTION

3.1 Vehicle Defects

To demonstrate the effectiveness of the proposed method, we use publicly available data from NHTSA, namely the Technical Service Bulletins (TSBS) and Recalls Datasets, which are available at this web site: <https://www-odi.nhtsa.dot.gov/downloads/> These datasets are representative of defect datasets commonly used by NHTSA analysts and share the following characteristics:

- An identifier (Bulletin Number or Record ID): A code associated with the defect report, which is not necessarily unique. In some cases, technicians will use the same code for similar, but non-identical defects. For SmarTxT, we use the Bulletin Number as the identifier for technical reports in the TSBS Dataset. For the Recalls Dataset, the Record ID is used as the identifier.
- A set of metadata (NHTSA ID, Bulletin Date, Component Name, Make, Model, Year, Date Added): A set of values, generally categorical, that provide an outline of the problem that was identified. The NHTSA ID is an ID provided by NHTSA for each new technical report.

The Bulletin Data denotes when the original bulletin was drafted. The Component Name provides the code and description of the failing component. The Make, Model, and Year are used to identify a specific car (e.g., Ford Focus 2012 or Mazda CX-5 2021). The Date Added is the date when the bulletin was added into the NHTSA database.

- A textual description (Technical Summary): A description of the problem or other information associated with the defect report (e.g., in some cases there is also a description of the proposed fix).

An example defect report from the TSBS dataset is shown below in Figure 1.

Bulletin Number	NHTSA ID	Bulletin Date	Component Name	Make	Model	Year	Date Added	Technical Summary	
0	SSM 46303	10095164	20170113	060000 ENGINE (PWS)	FORD	FOCUS	2014	20170315	THE 2012-2017 FOCUS WORKSHOP MANUAL (WSM) PROCEDURE FOR OIL PAN AND REAR MAIN OIL SEAL INSTALLATION HAS BEEN REVISED. REFER TO WSM, SECTION 303-01 REMOVAL AND INSTALLATION. THE SEALANT APPLICATION LOCATION HAS BEEN REVISED. SPECIAL ATTENTIO
1	SSM 46312	10095165	20170118	110000 ELECTRICAL SYSTEM	FORD	ESCAPE	2017	20170315	SOME 2017 ESCAPE VEHICLES EQUIPPED WITH SYNC CONNECT AND ORDERED WITH A FORD ORIGINAL ACCESSORY(FOA)/FACTORY INVOICED ACCESSORY (FIA)REMOTE START SYSTEM MAY EXHIBIT A CONCERN THAT THE ENGINE CANNOT BE REMOTE STARTED USING THE FORDPASS APPLI
2	PI1276A	10095166	20140910	102000 POWER TRAIN:MANUAL TRANSMISSION	CHEVROLET	CORVETTE	2014	20170315	THIS PRELIMINARY BULLETIN PROVIDES A PROCEDURE TO ADJUST THE SHIFTER TO CORRECT THE CUSTOMERS CONCERN OF THE TRANSMISSION HARD TO SHIFT OR POPS OUT OF GEAR.
3	PI1276A	10072502	20140901	102000 POWER TRAIN:MANUAL TRANSMISSION	CHEVROLET	CORVETTE	2014	20160622	SUMMARY TO BE PROVIDED ON A FUTURE DATE.
4	SSM 46317	10095167	20170119	353400 EQUIPMENT:ELECTRICAL:RADIO/TAPE DECK/CD ETC.	FORD	F-150	2017	20170315	SOME 2017 F-150 VEHICLES EQUIPPED WITH SATELLITE RADIO MAY EXHIBIT AN ACQUIRING SIGNAL MESSAGE DISPLAYED FOR GREATER THAN 5 MINUTES AND NO SOUND FROM THE SPEAKERS WHEN SATELLITE RADIO IS SELECTED. THIS CONCERN OCCURS WHEN THERE ARE NO OBSTR

Figure 1: Example TSBS Defect Reports. We use the Bulletin Number as the record identifier and apply our Natural Language Processing workflow to the Technical Summary.

As shown above, in Figure 1, the Technical Summary provides the textual description for the TSBS dataset. The TSBS dataset downloaded directly from NHTSA, contains almost 3M defect reports at the time of writing, with new reports being added daily. However, we find many duplicated reports, namely, we find multiple reports where the BULNO (i.e., the identifier) and the Summary match exactly. We remove duplicated records but add a column to note the number of duplicates to provide analysts with an idea of the frequency of each defect. After removing duplicated records, the dataset is reduced to 160k defect reports. The Recalls dataset contains 190k recall reports with no duplicates, so no reports are removed from this dataset. Once the duplicated records are removed, we then proceed to preprocess the text fields to be appropriate for applying NLP.

The technical defect reports that we analyze with NLP techniques are often written as a memo for the technicians rather than explicit documentation of the issue. This means that the language used is technical and the specific issue may not be clearly stated. Our selected class of neural networks, called transformers, can link words with the context and link sentences with other sentences, which helps to address this kind of language problem.

3.2 Text Preprocessing for NLP

NLP is generally applicable for a wide range of text-based datasets; however, it is generally recommended to perform a set of preprocessing steps to improve the inference driven from NLP

models [31]. Our text preprocessing procedure is completely automated and follows the six steps outlined below:

1. Remove punctuation
2. Remove numbers
3. Tokenize the text
4. Remove stopwords
5. Lemmatize the text
6. Remove words less than three characters

For many NLP tasks, the punctuation and numbers contained in the text descriptions do not add meaningful information to understanding the text. This is the case for our study, where numbers are generally used for additional information and do not directly impact the defect report. One exception is the year range, which is often mentioned in the summary. However, the year is captured in the metadata of each defect report, so there is no loss of information. Similarly, punctuation, such as “.”, “;”, and “-” do not provide meaningful textual information for our study. Next, we tokenize the text and remove any stopwords. Stopwords (e.g., a, an, the) are context-independent words in English that do not impact the context of the text. We use a stopwords list provided by NLTK [32], a commonly used package for NLP. Next, we lemmatize the text using the WordNet Lemmatizer from NLTK [32]. The WordNet Lemmatizer replaces words with a base form of the word only if the base form is a valid dictionary word. For example, the words “removes” and “removed” would be lemmatized to “remove”. Similarly, the words “saving” and “saved” would be lemmatized to “save”. Lemmatization helps NLP models to learn the importance of individual words more quickly. Finally, we remove any words less than three characters, which are generally acronyms.

4. NLP APPROACHES: TOPIC MODELING AND DOCUMENT SIMILARITY

4.1 Topic Modeling

The objective of topic modeling is to identify latent semantic patterns in unstructured text that uncover previously unknown groupings in the text that might correspond to common vehicle issues. In topic modeling documents of unstructured text are clustered semantically based on words that appear in documents. The key assumption is that documents containing similar words belong to the same topic and describe similar problems. These topics are not known ahead of time and thus require unsupervised, probabilistic approaches, such as Latent Dirichlet Allocation (LDA) or Probabilistic Latent Sentiment Analysis (PLSA), where documents are modeled as a random mixture of topics and topics are characterized as a distribution of words. In these models, a “bag-of-words” approach is used, where word frequency is retained but the word order is lost. Because of this approach, word co-occurrence is strongly influential in determining the latent topics.

4.1.1 LDA Model

To identify topics in the NHTSA datasets, we use LDA, a generative probabilistic model, which was first introduced by Blei et al. in 2003 [16], where they provide the following definitions:

- A word (w) is the basic unit of data, which is an item from the vocabulary.

- A document (d) is a sequence of N words denoted by $d = (w_1, w_2, \dots, w_n)$, where w_n is the n th word in the sequence. In our case, a document is a pre-processed summary of a technical bulletin or a recall description.
- A corpus (D) is a collection of M documents denoted by $D = (d_1, d_2, \dots, d_n)$, where d_n is the n th document in the corpus. In our case, a corpus is a collection of technical bulletins or recalls.

Documents are represented as a random mixture over latent topics (Z), where each topic (z_i) is characterized by a distribution over words. LDA assumes a generative process for creating documents as presented in [33]:

1. Choose $\theta_i \sim Dir(\alpha)$
2. Choose $\phi_i \sim Dir(\beta)$
3. Choose topic $z_i \sim Multinomial(\theta_i)$
4. Choose word $w_i \sim Multinomial(\phi_{z_i})$

Here, α and β are hyperparameters for the document-topic (θ) and topic-word (ϕ) Dirichlet distributions, respectively. While these hyperparameters are typically symmetric, they are not required to be. LDA, while unsupervised, can be optimized through hyperparameter tuning and a model selection criterion, such as AIC or BIC, by comparing the log-likelihood of various models. This allows for an informed selection of the number of topics, alpha, and beta.

4.2 Document Similarity

4.2.1 Introduction to Document Embedding and Cosine Similarity

Document similarity is an NLP method that quantifies pairwise semantic measures of all documents for a given text field. Semantically similar documents can be queried for a reference document by sorting on other documents' similarity scores. Identifying semantically similar documents is particularly useful when locating other records related to a central problem summarized in a previous document [34]. In the TSBS dataset, document similarity is utilized to link similar detect defects in vehicles to another.

Semantic similarity is quantified by applying cosine distance to embedded document vectors. Embedded vectors represent the optimized conditional probability of a given word occurring in relation to all other words in the corpus [13] Capturing the conditional probability of each word w given all context words c can be represented by the following *softmax* function:

$$P(w | c) = \frac{\exp((u_w)^T v_c)}{\sum_{w \in Vocab} \exp((u_w)^T v_c)}$$

- Where the numerator captures the similarity between the input word u_w and the context word vector transposed from the weight matrix v_c
- The denominator provides a sum of all context words v_c the dot product of the input word u_w to normalize similarity for a given context word in relation to all other words.

In place of a simple embedding, we use a document embedding model which adds a paragraph id as another input vector. The product outputs from both the word embedding layer and the document embedding layer are concatenated together and averaged before applying the *softmax* activation function to produce the document embedded vector [14]. After generating these document vectors, we apply a cosine similarity metric to each document vector allowing for each document's semantic value to be quantified.

4.2.2 Bidirectional Encoder Representation from Transformers (BERT) document embeddings

BERT, a pre-trained document embedding model developed by Google in 2018 [18], was used to ensure validity when quantifying semantic values in technical summaries to capture document similarity. BERT performs superior to traditional embedding approaches due to BERT's ability to bidirectionally learn contextual relationships among phrases. This means that words that are positioned before or after a target word are considered for contextual relevance. This is unique from previous approaches, which only use one-directional embedding models to predict the next word in a sequence. Furthermore, LSTM models could only predict words immediately before and immediately after a given word [35]. In contrast, BERT predicts the contexts of word phrases as they relate to each word for a given text string. It does this through a Masked Language Modeling (MLM) approach that involves an encoder-decoder structure in which each encoder layer is fed into all decoder layers allowing for all input words to be processed simultaneously [18]. As a result, different contexts of the same word are captured which was not possible in prior embedding techniques. For instance, the word "gauge" could be distinguished between the noun form in "tire pressure gauge" from the verb form "gauge".

Even more so, BERT implements a Transformer model architecture. Transformer models apply constant steps of "attention" mechanisms that enhance the inference among connections between all words in a sentence. These attention mechanisms work by capturing both the hierarchical structure of language (parsing) along with context (composition) [18, 36]. The BERT model will quickly identify, or selectively attend to words in a sentence that are salient to the context of that sentence rather than identifying the word's overall context. More specifically, the attention mechanisms allow for local attention for a given word as it focuses on words of relative importance [36].

The outcome of the attention compatibility measures reflects aspects of syntax, or function to parse language. Concatenation of attention mechanisms into a multi-head within each layer capture context dictated by the individual word's local attention within each sentence. Coreference links are identified, or expressions that are semantically equivalent to phrases in other documents independent of syntax structure, through the stacking of multiple attention layers in the BERT architecture [37]. The pre-trained BERT base model was used for developing the initial SmarTxT prototype. The base model consists of 12 layers (transformer blocks), 12 attention heads, and 110 million parameters. The outcome of the iterative encoding-decoding attention mechanisms of the BERT model generates a dense vector representation of documents in which coreference links are

quantified. BERT's sophisticated embedding technique was used to generate document embeddings for the technical summaries to calculate similarity.

4.2.3 Fine-tuning BERT

We fine-tuned the BERT base model with the TSBS dataset to establish more data-specific document embeddings for capturing document similarity. Fine-tuning, also called transfer learning, is a common procedure in deep learning for large neural networks [38, 39, 40]. The motivating idea is to take a well-trained model in a similar domain and then retrain only the top layer(s) for a new domain. The model retains the general understanding of the domain, while also improving for a new domain through fine-tuning. The fine-tuning of document embeddings requires a supervised model architecture for the learned embeddings to be optimized. Therefore, pairwise documents are fed into the pre-trained BERT model with associated similarity scores for each pair. Reimers and Gurevych [41] demonstrate a fine-tuned sentence embedding approach using predefined similarity labels for each document pair. In this use case, similarity scores were captured by calculating cosine similarity for each document pair after embeddings were generated by the Doc2Vec model. Following this approach, a dictionary consisting of each document pair and associated similarity score was the input into the BERT model. The fine-tuning occurs through feeding each tokenized sentence into the BERT network followed by a mean-pooling layer. A pooling layer is a method for down-sampling and is the default for Siamese BERT-Networks, like the ones provided by the Python package *sentence-transformers* [41]. Siamese BERT-Networks are one type of fine-tuning scheme for BERT and the one we use. The averaged vector (i.e., the output of the mean-pooling layer) represents the pre-trained BERT vectors tuned on incoming sentence pairs and their respective similarity scores. Cosine similarity is calculated for the averaged vectors then compared against input similarity scores. Model weights and document embeddings are updated based on loss between two similarity scores. Tuning the pre-training model on the supervised output has been determined to make moderate changes to the attention mechanisms in the network. More specifically, fine-tuning BERT has shown to be most effective by enhancing domain-specific (in our case, vehicle defects) representations of sentences. In contrast, the pre-trained model is well adapted to identify global syntactic patterns found in everyday language.

5. SMARTXT APPLICATION

5.1 General Approach

The intended user of our application is a vehicle defect investigator at NHTSA or a similar role in another organization. A key challenge for current NHTSA vehicle defect investigators is the inability to access sensor data from vehicles directly. Instead, to perform defect investigations they must rely on technical bulletins and identify patterns and trends in the text data. Our tool, SmartXT provides NHTSA vehicle defect investigators the ability to identify key topics in the technical summaries and identify similarities among defects to expedite the in-depth vehicle investigations.

Our application provides 5 web pages to assist the user in investigating defects in the NHTSA data. A summary of each step is given below:

1. **Login:** A user logs in to the application and the application assigns the correct privileges (e.g., datasets that are visible to the specific user)

2. **Overview of the Dataset:** The user selects a dataset to analyze from the currently available datasets in the tool. A user can view the selected dataset contents and retrieve summary statistics.
3. **Train Models Page:** The user filters the dataset by make, model, and year (if needed) and trains the topic model on the filtered dataset or chooses to use the pre-trained NLP models for topic modeling and document similarity. The user can also see summary statistics for the filtered dataset.
4. **Topic Modeling:** The user views the results of topic modeling, where each topic along with the reports that best fit in each topic are displayed as shown in Figures 3a and 3b.
5. **Document Similarity:** The user provides the identifier of a specific document (the Bulletin ID in the case of Technical Bulletins) and the tool returns a list of documents (Figure 3) that have similar content, ranked in order of similarity

Topic Modeling

Topic	Topic Keywords	View Similar Documents
1	provided date summary future stock newsletter overview solutions pending campaigns	Topic 1
2	oil fuel model engine leak pump coolant valve new models	Topic 2
3	engine vehicles exhibit model equipped start condition year idle cold	Topic 3
4	vehicles campaign service certain warranty dealer parts vehicle recall program	Topic 4
5	bulletin service information technical provides procedure instructions update repair new	Topic 5
6	regarding subject information seat vehicles switch tire harness models belt	Topic 6
7	control module diagnostic dtc code trouble software lamp engine indicator	Topic 7

Figure 2a: SmarTxT Topic Modeling Results Page. Clicking on “Topic 1” generates

Documents with highest likelihood to belong in Topic 1

Topic 1 keywords: provided date summary future stock newsletter overview solutions pending campaigns

	Bulletin Number	Make	Model	Year	Date Added	Technical Summary	Topic_Likelihood
0	14941	FCA	ALL MODELS	2006	20171013	Headlamp Assy -Frt R/L You have been identified as either having stock, or have stock in transit to your dealership. Please return suspect stock of part numbers:68214396AB, & 68214397AB. See attachments for identifying suspect stock. Ple	0.880607
1	TT-ALI	SUBARU	SUBARU	9999	20180611	Tech Tips Article Locator Index	0.849986

. Figure 2b: SmarTxT Topic Modeling Results Deep Dive Page.

Reference Document

Bulletin Number	NHTSA ID	Bulletin Date	Component Name	Make	Model	Year	Date Added	Technical Summary	Best_Topic	
143039	SSM 48551	10171954	20200214	110000 ELECTRICAL SYSTEM	FORD	TRANSIT CONNECT	2020	20200303	SSM 48551 - 2020 Various Vehicles Equipped With SYNC 3 Various SYNC Performance Related Concerns Some 2020 Ford and Lincoln vehicles equipped with SYNC 3 may exhibit various performance related concerns such as: intermittent frozen or blan	6

Similar Documents

Bulletin Number	NHTSA ID	Bulletin Date	Component Name	Make	Model	Year	Date Added	Technical Summary	Best_Topic	
148227	SSM 48833	10176119	20200522	110000 ELECTRICAL SYSTEM	FORD	ESCAPE	2020	20200604	Some 2020 Ford and Lincoln vehicles may exhibit a symptom of no ability to schedule a remote start event using the FordPass or LincolnWay app. If only the scheduling portion of the remote start system is not functioning, please ensure that	6
31877	44872	10057837	20150109	113000 ELECTRICAL SYSTEM-STARTER ASSEMBLY	FORD	MUSTANG	2015	20150701	FORD: VARIOUS 2015 FORD MUSTANG VEHICLES EQUIPPED WITH THE REMOTE START FEATURE MAY EXPERIENCE DTC P025A STORED IN PCM MEMORY WITH CRANK NO-START OR STALL PROBLEMS. *TA	5
143256	SSM 48558	10172019	20200218	110000 ELECTRICAL SYSTEM	FORD	F-250 SD	2020	20200303	SSM 48558 - 2017-2020 Various Vehicles Equipped With Telematics Control Unit (TCU) - Various FordPass/LincolnWay Mobile App Concerns Some 2017-2020 Ford and Lincoln vehicles equipped with a TCU may exhibit inoperative remote features via F	6
122013	SSM 47444	10144930	20180802	110000 ELECTRICAL SYSTEM	FORD	EXPLORER	2016	20180924	Some 2016-2018 Ford and Lincoln vehicles equipped with SYNC 3 not at the latest software level may exhibit Various Performance Related concerns such as Display Operation, Navigation, Voice Recognition, Phone Pairing, AppLink, Travel Link, S	6
109733	SSM 46860	10125078	20171006	353500 EQUIPMENT-ELECTRICAL-NAVIGATIONAL SYSTEM(GLOBAL POSITIONING SYSTEM)	FORD	ESCAPE	2016	20171214	Some 2016-2018 Ford or Lincoln vehicles equipped with SYNC 3 may exhibit a blank screen or other various SYNC symptoms only when the vehicle is in transport mode. This is due to the power saving features enabled while in transport mode and	6
109734	SSM 46861	10125079	20171006	353000 EQUIPMENT-ELECTRICAL	FORD	ESCAPE	2016	20171214	Some 2016-2017 Ford and Lincoln vehicles with SYNC 3 may intermittently display an overhead image of a vehicle in the center of the SYNC touchscreen with no obstacle detection grids or warning chimes. SYNC may appear unresponsive when the o	7

Figure 3: SmarTxT Document Similarity Results

5.2 SmarTxT Implementation

Our application was built using Python, Flask, and Bootstrap. Flask was selected as the web framework of choice due to its ease of use with Python and customization features. The web pages were developed using Bootstrap to match current industry standards and allow for the reuse of code throughout various parts of the application. All other code, including the NLP models, was written in Python. All timing data was calculated on a Dell Precision 5540 with an i7-9850H CPU and 32GB of RAM. Key aspects of the implementation are provided below in further detail:

- **Multiple dataset compatibility:** To increase the effectiveness of the tool, we built the functionality to analyze distinct datasets. To achieve this, Python code was written functionally, such that specific dataset parameters were passed throughout the app instead of hard-coded. For each dataset, a JSON file is created with the appropriate metadata, which is then processed once a user selects a dataset to analyze.
- **Topic Modeling:** Topic modeling (i.e., LDA) was implemented using scikit-learn in Python [42] to provide results to users in real-time. The user selects the number of topics, but all other hyperparameters (i.e., alpha, beta, and the learning decay) are set to the default values provided by scikit-learn. From the topic modeling implementation in scikit-learn, we provide the user information on the topics, as well as the reports that best fit into each topic to identify key defect reports more quickly in the dataset. For small datasets, topic modeling is computationally fast (i.e., less than 10 seconds). Thus, when users select any filters (i.e., make, model, year) in SmarTxT a topic model is trained at runtime based on the filters selected. However, with the large size of the NHTSA datasets (e.g., 180k+ records) topic modeling was too slow to apply at runtime, when analyzing the entire dataset. To avoid the user experiencing this slowdown, we implemented an offline pre-

training procedure where the topic model for the entire dataset is pre-computed. For this model, we perform `RandomSearchCV` in `scikit-learn` to identify the optimal topic model based on the log-likelihood. We test the following hyperparameters and ranges: number of topics (3, 10), alpha (0, 1), beta (0, 0.5), and learning decay (0.5, 0.99). Our procedure generates 90 candidate topic models and returns the best model based on log-likelihood. Then the topic model is loaded into `SmarTxT` if a user does not filter the dataset. This option is important for analysts who want to investigate defects that may affect multiple makes or models.

- **Document Similarity:** Document Similarity was implemented using the `Sentence Transformers` package in Python, which provides sentence-level embeddings using BERT [41]. A user will provide a reference document and all similar documents are returned to the user. We only use the text field (e.g., technical summary for TSBS) to identify similar documents. As mentioned previously, cosine similarity was the metric of choice for identifying similarity between documents. We return any document that has a cosine similarity score greater than a threshold value. To identify the similarity threshold to use, we tested multiple possible thresholds (e.g., 0.7, 0.8, and 0.9) for each dataset. Due to the lack of a validation set, we are unable to optimize the similarity threshold. Instead, we select the similarity threshold based on the average number of similar documents returned on average for each dataset. For the TSBS dataset, the selected similarity threshold was 0.8, while the selected similarity threshold for the Recalls dataset was 0.9. The difference in similarity thresholds can be explained by the lack of exact duplicates in the Recalls dataset. Instead, there are many similar, but non-identical recalls. This is dissimilar to the TSBS dataset, where many documents are exact duplicates, which we remove in our pre-processing procedure. Similar to topic modeling, performing document similarity at runtime is too slow for a user due to the large amount of computation required to calculate the BERT embeddings and perform the similarity calculation. To avoid the user experiencing this slowdown, we perform a global similarity calculation offline and save the result into a lookup table, which is described further in Section 4.3.

`SmarTxT` is deployed in the AWS cloud using `CloudFront`, `EC2`, and `Lambda`. All cloud operations complete in less than 15 seconds, which is the timeout applied for our `Lambda` functions. We use `CloudFront` to deploy the application with code stored in `S3` and `EC2`. Specifically, we use an AWS Linux `t2.large` `EC2` instance. Routes within the app call a `Lambda` function that invokes the function on `EC2` and then returns the appropriate data to the webpage.

5.3 Custom Implementation to Improve User Experience

A common tactic for deploying machine learning models in production is to pretrain models offline and then make predictions at runtime. This improves the responsiveness and performance of the application, without sacrificing model accuracy.

Document similarity requires computing the embedding of each Technical Summary and then computing the cosine similarity between every combination of Technical Summaries. This process is quite slow, so we perform these steps offline. In general, to compute a similarity matrix for a vector v with n entries would create a matrix A , where A is $N \times N$. For the TSBS dataset, we have 157k entries, which would lead to a matrix requiring well over 100 GB to store in memory. Instead,

we compute the similarity matrix in batches, where batch size $b < n$. This leads to a smaller matrix B , where B is $b \times N$. We select b such that B will fit in the available RAM. Once B is constructed, each row i denotes the document in the i th index of the TSBS dataset batch ($i \in \{1, 2, \dots, b\}$) and each column j denotes the document in the j th index of the TSBS dataset ($j \in \{1, 2, \dots, n\}$). $B_{i,j}$ captures the similarity between document i and document j . For each document i , we first remove any entries with a similarity score less than the threshold. If more than 100 documents have a similarity score above the threshold, we keep the 100 most similar documents. The 100-document limit is selected as a cutoff to balance 1) the size of the lookup table being generated and 2) to avoid overwhelming NHTSA analysts with too much information. As the tool progresses, we plan to work with NHTSA to identify the optimal number of similar documents to return for each query. This approach provides analysts with a selection of most similar documents, without leading to an overwhelming amount of data. After applying this procedure for each document and each batch, the indices j of the most similar documents are saved to a lookup table. When a user enters a reference document, SmarTxT queries the lookup table to provide the user with the most similar documents. We find that it takes 64.2 minutes on average to generate the lookup table for the TSBS data and 68.76 minutes to generate the lookup table for the Recalls dataset. The discrepancy in time can be attributed to the additional records in the Recalls dataset (roughly 30k more). It is important to note that if a user applies any of the optional filters (make, model, year) prior to the document similarity page, these filters are applied after querying the lookup table. Thus, SmarTxT will always display the most similar documents, according to BERT, for any query regardless of the filtering options selected by the user. The query is nearly instantaneous and provides the user with a significantly faster and more usable experience. This implementation provides the user a faster, more responsive experience with our application, which in turn could lead to faster adoption due to ease of use.

6. RESULT VALIDATION

As previously mentioned, document similarity is an unsupervised task, which means that we do not have a set of labels to validate the results we have obtained. Thus, we use a human expert to validate our results. While this limits the amount of validation we can perform, due to the time required to manually review each defect report, it provides a repeatable benchmarking approach that accounts for the subjectivity of the outcome. We have randomly selected 10 documents and for each of these documents, we have identified the top 10 similar documents using our approach. Then we have submitted the 100 resulting documents to a domain expert and asked the expert to review the results and give a subjective rating of their similarity to the reference document. The instructions were as follows:

1. Read the reference document.
2. Read each document deemed similar by our approach and for each document:
 - Rate it 0 if in the expert opinion the document does not describe a problem that is similar to the one described in the reference document.
 - Rate it 1 if in the expert opinion the document describes a problem that is similar to the one described in the reference document.

At the end of this process, the human expert had reviewed 100 documents and obtained the ratings shown in Table 1. From these observations, we see that our approach to document similarity on

this specific dataset gives us 83% accuracy. The observations we made are available for review upon request. In general, we can say that our document similarity method routinely provided relevant documents that would be beneficial for NHTSA or other analysts.

Table 1: Validation Results for 10 Randomly Selected Documents

Reference Document	Not Similar	Similar
TSB 18-2352	2	8
99411	3	7
12925	1	9
SSM 46528	1	9
SSM 49061	3	7
04L23	2	8
996105	2	8
976011	1	9
TSB 20-2100	0	10
15-0109	2	8
Sum >>>>>>		83

7. CONCLUSIONS

In conclusion, we have shown that natural language processing can be used effectively to extract patterns from a dataset documenting vehicle defects. SmarTxT possesses multiple advantages for NHTSA analysts and consumers of the NHTSA datasets: it's intuitive, it has built-in models, and provides a way to explore multiple datasets.

Currently, we are experimenting with more sophisticated approaches to elicit topics by extracting features. In the last years, there has been growing interest in topological data analysis as a tool for extracting features out of a textual dataset [43]. This is a new development because, normally, topological techniques are not used for natural language processing. We have experimented with the Mapper algorithm in a different dataset from the one used in this paper and have obtained satisfactory results. Additionally, we are developing automated administrative functionality to allow users to add datasets to the tool directly. A user would provide a dataset, which would then be processed, pretrained with topic modeling and document similarity. Further research will be needed to automatically identify the optimal similarity threshold for a new dataset. However, we envision this will require a validation set, which we did not have in this study. Automating these steps would allow for users to upload any dataset to be used in SmarTxT. Additionally, an interesting research avenue is to determine if adding metadata would improve the quality of the returned similar documents. For example, using the make, model, and year columns along with the text could improve the quality. The key challenge in this approach is identifying the correct weighting scheme for each field and handling highly correlated fields (such as the make and model which are often discussed in the text. Further research is needed to determine the relationship between the structure of the models used and the data analyzed. For example, we used Siamese-BERT for fine-tuning while other options for finetuning exist in the literature [41]. We plan to

explore the effect of model structure, additional hyperparameter tuning, and training BERT from scratch using vehicle defect data to develop a high-quality model for NHTSA and other organizations to use in the investigation of vehicle defects. Further, these research efforts will provide a better understanding of how to extend the use of a similar application to other fields of study, such as LEGAL-BERT [23] and BioBERT [22] that have been tuned for the legal and medical industries, respectively.

10. REFERENCES

- [1] T. R. B. A. I. Committee, "A Primer on Machine Learning for Transportation," in *Transportation Research Board Annual Meeting*, 2020.
- [2] C. Wu, A. Kreidieh, K. Parvate, E. Vinitzky and A. M. Bayen, "Flow: Architecture and Benchmarking for Reinforcement Learning in Traffic Control," *CoRR*, vol. abs/1710.05465, 2017.
- [3] J. Booth, B. Di Eugenio, I. F. Cruz and O. Wolfson, "Robust Natural Language Processing for Urban Trip Planning," *Applied Artificial Intelligence*, vol. 29, no. 9, pp. 859-903, 2015.
- [4] L. Barua, B. Zou and Y. Zhou, "Machine Learning for International Freight Transportation Management: A Comprehensive Review," *Research in Transportation Business & Management*, vol. 34, pp. 1004-53, 2020.
- [5] A. P. Sligar, "Machine Learning-Based Radar Perception for Autonomous Vehicles Using Full Physics Simulation," *IEEE Access*, vol. 8, pp. 51470-6, 2020.
- [6] A. Ata, M. A. Khan, S. Abbas, M. S. Khan and G. Ahmad, "Adaptive IoT Empowered Smart Road Traffic Congestion Control System Using Supervised Machine Learning Algorithm," *The Computer Journal*, 2020.
- [7] J. Yang, S. Li, Z. Wang, H. Dong, J. Wang and S. Tang, "Using Deep Learning to Detect Defects in Manufacturing: A Comprehensive Survey and Current Challenges," *Materials*, vol. 13, p. 5755, 2020.
- [8] W. Liqun, W. Jiansheng and W. Dingjin, "Research on Vehicle Parts Defect Detection Based on Deep Learning," *Journal of Physics: Conference Series*, vol. 1437, no. 1, 2020.
- [9] P. M. Nadkarni, L. Ohno-Machado and W. W. Chapman, "Natural Language Processing: an Introduction," *Journal of the American Medical Informatics Association*, vol. 18, no. 5, pp. 544-551, 2011.
- [10] NHTSA, "Risk-Based Processes for Safety Defect Analysis and Management of Recalls," NHTSA, 2020.
- [11] NHTSA, "General Standard Operating Procedures (SOP)," NHTSA, 2017.

- [12] S. Russell and P. Norvig, "Chapter 23: Natural Language Processing," in *Artificial Intelligence: A Modern Approach 4th Edition*, Pearson Education, 2020, pp. 823-855.
- [13] T. Mikolov, K. Chen, G. Corrado and J. Dean, "Efficient Estimation of Word Representations in Vector Space," *arXiv preprint arXiv:1301.3781*, 2013.
- [14] Q. Le and T. Mikolov, "Distributed Representations of Sentences and Documents," in *International Conference on Machine Learning*, 2014.
- [15] J. Ramos, "Using TF-IDF to Determine Word Relevance in Document Queries," in *First Instructional Conference on Machine Learning*, 2003.
- [16] D. M. Blei, A. Y. Ng and M. I. Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research*, vol. 3, p. 993–1022, 2003.
- [17] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. a. L. R. Rault, M. Funtowicz and others, "Huggingface's Transformers: State-of-the-art Natural Language Processing," *arXiv preprint arXiv:1910.03771*, 2019.
- [18] J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, "Bert: Pre-training of Deep Bidirectional Transformers for Language Understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [19] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee and L. Zettlemoyer, "Deep Contextualized Word Representations," *arXiv preprint arXiv:1802.05365*, 2018.
- [20] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell and others, "Language Models are Few-Shot Learners," *arXiv preprint arXiv:2005.14165*, 2020.
- [21] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan and S. Gelly, "Parameter-Efficient Transfer Learning for NLP," in *International Conference on Machine Learning*, 2019.
- [22] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So and J. Kang, "BioBERT: a Pre-trained Biomedical Language Representation Model for Biomedical Text Mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234-1240, 2020.
- [23] I. Chalkidis, M. Fergadiotis, P. Malakasiotis, N. Aletras and I. Androutsopoulos, "LEGAL-BERT: The Muppets Straight out of Law School," *arXiv preprint arXiv:2010.02559*, 2020.
- [24] F. Gurcan and N. E. Cagiltay, "Big Data Software Engineering: Analysis of Knowledge Domains and Skill Sets using LDA-based Topic Modeling," *IEEE Access*, vol. 7, pp. 82541-82552, 2019.

- [25] H. Jelodar, Y. Wang, C. Yuan, X. Feng, X. Jiang, Y. Li and L. Zhao, "Latent Dirichlet Allocation (LDA) and Topic Modeling: Models, Applications, a Survey," *Multimedia Tools and Applications*, vol. 78, no. 11, pp. 15169-15211, 2019.
- [26] L. Liu, L. Tang, W. Dong, S. Yao and W. Zhou, "An Overview of Topic Modeling and its Current Applications in Bioinformatics," *SpringerPlus*, vol. 5, no. 1, pp. 1-22, 2016.
- [27] H. Alambeigi, A. D. McDonald and S. R. Tankasala, "Crash Themes in Automated Vehicles: A Topic Modeling Analysis of the California Department of Motor Vehicles Automated Vehicle Crash Database," *arXiv preprint arXiv:2001.11087*, 2020.
- [28] T. H. Haveliwala, A. Gionis, D. Klein and P. Indyk, "Evaluating Strategies for Similarity Search on the Web," in *11th International Conference on World Wide Web*, 2002.
- [29] F. Akhbardeh, T. Desell and M. Zampieri, "NLP Tools for Predictive Maintenance Records in MaintNet," in *1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: System Demonstrations*, 2020.
- [30] R. McEwan, G. B. Melton, B. C. Knoll, Y. Wang, G. Hultman, J. L. Dale, T. Meyer and S. V. Pakhomov, "NLP-PIER: a Scalable Natural Language Processing, Indexing, and Searching Architecture for Clinical Notes," *AMIA Summits on Translational Science Proceedings*, vol. 2016, 2016.
- [31] S. Kannan, V. Gurusamy, S. Vijayarani, J. Ilamathi, M. Nithya, S. Kannan and V. Gurusamy, "Preprocessing Techniques for Text Mining," *International Journal of Computer Science & Communication Networks*, vol. 5, no. 1, pp. 7-16, 2014.
- [32] S. Bird, E. Loper and K. Ewan, *Natural Language Processing with Python*, O'Reilly Media Inc., 2009.
- [33] T. L. Griffiths and M. Steyvers, "Finding Scientific Topics," *Proceedings of the National Academy of Sciences*, vol. 101, p. 5228-5235, 2004.
- [34] M. Oghbaie and M. M. Zanjireh, "Pairwise Document Similarity Measure Based on Present Term Set," *Journal of Big Data*, vol. 5, p. 1-23, 2018.
- [35] V. Maslej-Krešňáková, M. Sarnovský, P. Butka and K. Machová, "Comparison of Deep Learning Models and Various Text Pre-Processing Techniques for the Toxic Comments Classification," *Applied Sciences*, vol. 10, p. 8631, 2020.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser and I. Polosukhin, "Attention is All You Need," in *Advances in Neural Information Processing Systems*, 2017.

- [37] K. Clark, U. Khandelwal, O. Levy and C. D. Manning, "What Does BERT Look At? An Analysis of BERT's Attention," *arXiv preprint arXiv:1906.04341*, 2019.
- [38] Y. Peng, S. Yan and Z. Lu, "Transfer Learning in Biomedical Natural Language Processing: an Evaluation of BERT and ELMo on Ten Benchmarking Datasets.," *arXiv preprint arXiv:1906.05474*, 2019.
- [39] M. Mozafari, R. Farahbakhsh and N. Crespi, "A BERT-based Transfer Learning Approach for Hate Speech Detection in Online Social Media.," in *International Conference on Complex Networks and Their Applications*, 2019.
- [40] C. Sun, Q. Xipeng, X. Yige and H. Xuanjing, "How to Fine-Tune BERT for Text Classification?," in *China National Conference on Chinese Computational Linguistics*, 2019.
- [41] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, 2019.
- [42] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825-2830, 2011.
- [43] S. Gholizadeh, K. Savle, A. Seyeditabari and W. Zadrozny, "Topological Data Analysis in Text Classification: Extracting Features with Additive Information," *arXiv preprint arXiv:2003.13138*, 2020.
- [44] C. Sun and Z. Yang, "Transfer Learning in Biomedical Named Entity Recognition: An Evaluation of BERT in the PharmaCoNER Task," in *5th Workshop on BioNLP Open Shared Tasks*, 2019.